

# Revealing Contentious Concepts Across Social Groups

Ching-Sheng Lin<sup>1</sup>, Zumrut Akcam<sup>1</sup>, Samira Shaikh<sup>1</sup>, Sharon Gower Small<sup>2</sup>, Ken Stahl<sup>1</sup>,  
Tomek Strzalkowski<sup>1</sup>, Nick Webb<sup>3</sup>

<sup>1</sup>ILS Institute, University at Albany, State University of New York

<sup>2</sup>Department of Computer Science, Siena College

<sup>3</sup>Department of Computer Science, Union College

{clin3, zakcam, sshaikh, kstahl, tomek}@albany.edu, ssmall@siena.edu, webbn@union.edu

## Abstract

In this paper, a computational model based on concept polarity is proposed to investigate the influence of communications across the ‘diacultural groups’. The hypothesis of this work is that there are communities or groups which can be characterized by a network of concepts and the corresponding valuations of those concepts that are agreed upon by the members of the community. We apply an existing research tool, ECO, to generate text representative of each community and create community specific Valuation Concept Networks (VCN). We then compare VCNs across the communities, to attempt to find ‘contentious concepts’, which could subsequently be the focus of further exploration as points of contention between the two communities. A prototype, CPAM (Changing Positions, Altering Minds), was implemented as a proof of concept for this approach. The experiment was conducted using blog data from pro-Palestinian and pro-Israeli communities. A potential application of this method and future work are discussed as well.

**Keywords:** diacultural groups, Valuation Concept Networks, contentious concepts

## 1. Introduction

Increasingly, we see communities expressing themselves online, through blogs or forums, to post opinions on specific issues. Groups of users congregate together around particular forums or blogs to express themselves in relation to new or recurring events. For our research, we wanted to use these posts as textual representations of a valuation system agreed upon by specific diacultural groups. In this study, we aim to discover the contentious concepts between opposing diacultural groups; to find those concepts which occur in both communities’ postings, but which are valued differently in each. At an elementary level, concepts can be thought as singularities (e.g., persons, locations, organizations, events) that are invoked by appropriate references in various forms of communication. We will identify these units, along with their valuation in the online communications of a community, and compare them with similar concepts in a different, opposing community. The potential of this technique is in use for government or business units to identify and monitor different points of view, with respect to specific issues, or across specific groups.

We created a prototype, CPAM (Changing Positions, Altering Minds), as a proof of concept of this approach. CPAM consists of two components: the first identifies concepts in text, and their associated valuations; the second explores the inter-conceptual relationships within the individual community, and then compares relationships between communities to locate contentious concepts.

The remainder of this paper is set out as follows. In Section 2, we review related research approaches, and discuss how they can be applied to our CPAM objective.

In Section 3, we describe the components of the proposed technique and the way they are used to establish CPAM. In Section 4, we discuss initial empirical studies, including data collection and evaluation. In final section, we present conclusions and future work.

## 2. Related work

The automatic detection of opinions and sentiment in text (cf. (Wiebe, Wilson and Cardie, 2005; Breck, Choi and Cardie, 2007; Strapparava and Mihalcea, 2008)) and speech (cf. (Vogt, Andre’ and Bee, 2008)) is a rapidly emerging area of research interest. Most initial work has focused on reviews of various kinds (Hu and Liu, 2004; David and Pinch, 2006), where users can be expected to be clear in signaling both positive and negative emotions in text. Most of the approaches use either a shallow, bag-of-words approach to the problem (calculating the number of ‘opinion’ words within some window around potential candidate judgment phrases, for example), or deeper, complex syntactic or dependency based approach to analysis. In common with prior work, we depend on a subjectivity lexicon (derived from the MPQA corpus (Wiebe, Wilson and Cardie, 2005)) and opt for a mechanism that is of a deeper level of understanding than bag-of-words, and yet does not necessitate deeper syntactic relationships. For CPAM, initial concept identification and annotation is based on our prior work in ECO (Effective Communication Online) (Small, Strzalkowski and Webb, 2010). The purpose of ECO is to extract and model the valuation system of the community and compare whether the contents of a new message fits into the targeted community. Users are guided in ways to shape their communication such that it eliminates, or mitigates the number of conflicting concept valuations between the new message and the concept representation of the target community. To achieve this, we must first

derive the salient concepts and corresponding polarities for the targeted community. We do this using a Transformation-Based Learning (TBL) approach, using lexical items, concepts, POS labels and the presence of polarity words in the input as learning features, and producing a Valuation System Vector (VSV). Accordingly, the same mechanism will be applied to the new message to obtain the Message System Vector (MSV). Finally, a comparison between two vectors is performed, identifying mismatches of concepts, and highlighting these to the user for the possible amendment. An example of this comparison of vectors between pro-Israeli message and pro-Palestinian blog is shown in Figure 1.

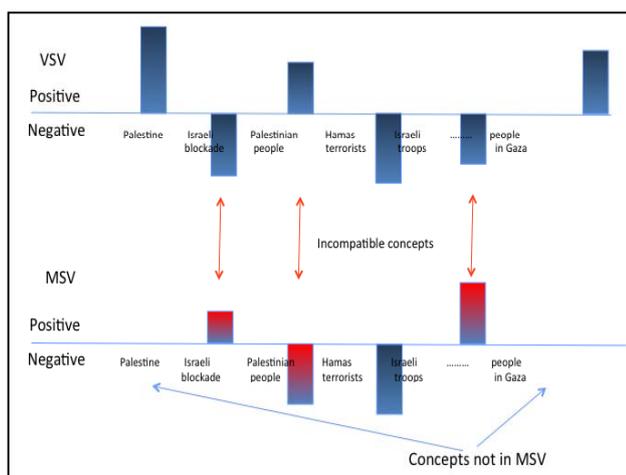


Figure 1: Comparing a MSV (bottom) from a pro-Israeli message to a VSV (top) from a pro-Palestinian blog. Incompatible concepts are highlighted in red.

### 3. CPAM methodology

The main goal of CPAM is to use the Valuation Concept Networks derived from community blogs to discover contentious concepts between two communities that are potentially open to negotiation. In this section, we describe the fundamental components for CPAM and explain how we use them to build the ultimate model for finding the contentious people concepts.

#### 3.1 Valuation Concept Network

Taking a Valuation System Vector from blog material belonging to a particular community, we extend that vector by including inter-conceptual temporal relationships for each community. By temporal relationships, we mean ‘temporal in text’, where one concept occurs prior to another in the source material, rather than existing in any formal temporal relationship. We believe that concepts as they occur in text are not accidental, that much like newspaper text, the occurrence of concepts is deliberate, and that one may be able to infer loose, causal relationships between two valued concepts. For example, if there are only two valued concepts in a blog posting, one is in the first sentence, and one is in the last sentence, these two concepts are still considered to be next to each another. These concepts are related to each

other, because they occur in the same discourse and because the focus (or activation) shifts from the first to the next, but not because there is any explicit textual or semantic relationship between them. Clearly, a semantic relationship may be postulated to present as well; however, extracting it with a degree of accuracy has eluded NLP research thus far because it requires detailed domain knowledge that is very expensive and time consuming to acquire. Our method bypasses this knowledge and domain dependency gridlock instead looking at the concept activation structure in a cultural narrative, irrespective of any specific subject domain. Take the example from an Israeli blog, detailing the barricade erected around Israeli settlements: “the barriers positive effect is in the control of arms to Palestinian terrorists”. Here ‘barrier’ is positively valued, and ‘Palestinian terrorists’ is negative, and we might infer that the positive value of the former is in controlling the negative concept that is the latter. For each community, we extract all the concepts and their valuations, then for each concept, we form a sequence of relationships by looking at the sequence of concepts in the text. Effective way is to create n-gram sequences of concepts based on their occurrence in source text per community.

After creating these sequences of relationships for each community, we use an open source graph visualization software, Graphviz, to display the resulting network structure. In these representations, node size represents the overall frequency of the concept and color represents the net cumulative count of negative and positive instances across all blog texts. Edge direction represents the temporal ordering of concepts mentioned in posting and edge thickness represents the frequency with which the relationship was noted. In Figure 2, the networks for pro-Palestinian and pro-Israeli blog data are shown.

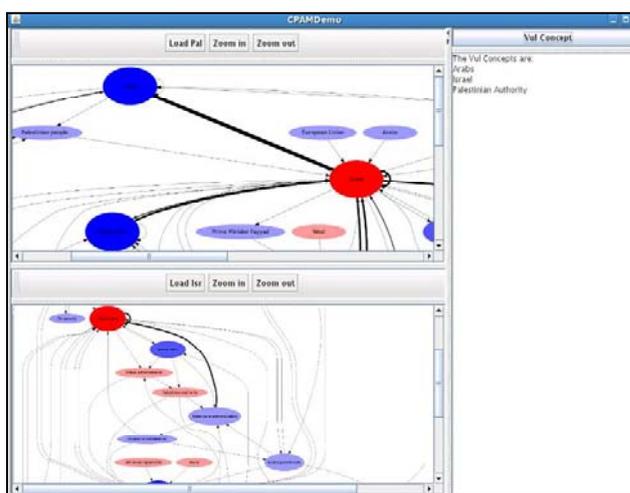


Figure 2. Left-top is the VCN for pro-Palestinian and left-bottom is for pro-Israeli.

#### 3.2 Contentious Concepts

The goal of CPAM is to find potential and display

contentious concepts. Our definition of a contentious concept is one that is shared between networks (therefore appears in blog postings of both communities) but the valuation of this concept in one network is the opposite of its valuation in the other network. Furthermore, we require that this concept is linked to other shared concept within the network where the valuation in both networks is the same. A contentious concept example between pro-Palestinian and pro-Israeli is show in Figure 3.

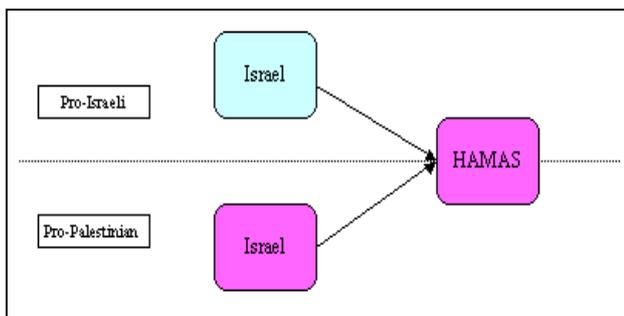


Figure 3. Two communities share their valuation of the concept “HAMAS” (negative). The two communities have different valuations of the concept “Israel”.

Once the networks of two communities are created, the next step is to identify the contentious concepts by comparing all relationships between two valuation concept networks.

#### 4. Data Collection

In this section, we will explain how we start with the initial set of experiments that eventually led to an extended set of experiments and discuss the results for generation of contentious concepts.

##### 4.1 Manual Annotation

The first step of data generation starts with manual annotation. English-language posting from two blog communities, one pro-Palestinian and one pro-Israeli, were collected for analysis and evaluation. We identified blog sites that are current and active, and are self-identified as representing the interests of their community. From those blog sites, we then have collected postings with the total of 36000 words for pro-Palestinian data and 27000 words for pro-Israeli data. Within these texts there are 577 and 362 concepts respectively that are assigned judgments by our multiple annotators. Of those, there are 335 and 213 unique concepts in the pro-Palestinian and pro-Israeli data sets. We then created the networks from these concepts as described in Section 3. These networks contain relationships of varying distance in text, so in order to limit data sparsity, we used the threshold distance of 2 in our initial experiments; this means we consider two concepts as related if there is at most one more concept mentioned between them. It should be noted that the actual textual distance (in words) between such concept mentions may be quite substantial. In the pro-Palestinian data there are 885 relationships

between concepts within window of 2 and 552 in the pro-Israeli data. Summary information of the initial data sets can be seen in Table 1. We then extract interesting, contentious concepts, as per our definition in Section 3. In the initial dataset, we found only 3 contentious concepts ("Israel", "Palestinian Authority" and "Arabs") that matched our definition. We can see from Table 1 that there are 335 unique concepts in pro-Palestinian and 213 in pro-Israeli. However, we calculated the number of shared concepts between the two data sets consisted of only 34 concepts in total. It became clear that this initial data set is not sufficient and we needed to expand our data set to obtain more significant contentious concepts.

	Pro-Palestinian Data	Pro-Israeli Data
Words	36,000	27,000
Concepts	577	362
Unique Concepts	335	213
Relationships	885	552

Table 1: Data summary of our initial data sets.

##### 4.2 Transformation-based Annotation

Given the summary data in Table 1, we estimated we needed four times as many words per community, to obtain around 2000 total concepts per data set. We collected an additional 200,000 words of data from each set of blog sites. Annotating this data manually was not feasible, so we used our initial data to train our ECO annotation tool, based around a transformation-based learning approach. This data was automatically processed, including removing HTML tags, part-of-speech tagging, AeroText concept spotting, and extraction of valuation words based on the MPQA corpus (Wiebe, Wilson and Cardie, 2005). The corpus was then processed through the ECO tagger to assign valuations to each concept. This involved using a set of rules acquired from the annotated data using transformation-based machine learning method, implemented under ECO. Applying the resulting automated tagger to our new data, we found around 2000 new concepts per community as predicted. Of those, we found that there were around 775 unique concepts per data set. This processed dataset becomes the input to the CPAM system. Summary information for the expanded data can be seen in Table 2.

	Pro-Palestinian Data	Pro-Israeli Data
Words	200,000	227,000
Concepts	2,177	2,676
Unique Concepts	773	775
Relationships	3,251	3,610

Table 2: Data summary for expanded, automatically annotated data set.

With the new data set generated, we derive concept networks from each data set, and then calculate the overlap relationships between them, and finally apply our criteria for selecting contentious concepts. We see that there are 202 overlapping concepts that appear in both data sets, which provides us a greater opportunity to find contentious concepts as per our criteria. Of those 202 concepts, we find 38 distinct contentious concepts, listed in Figure 4. These are concepts where the valuation of the concept is different between the two communities, yet they are connected to other concepts (not shown) which share valuations between communities. We noted that we can conflate some concepts ("Israel", "state of Israel"), using word overlap and synonymy, and that there were errors in the automatic processing ("U" is cut off from "U S A"). Nonetheless, it is interesting to see that well-known, highly contentious concept between these communities, such as the "Goldstone Report" are automatically identified.

MB, Obama, Egypt, Israel, Hamas, IDF, state of Israel, Israeli government, Middle East, West Bank, Goldstone, Gaza, Iran, Iraq, South Africa, Syria, U, US, PA, Gaza Strip, Jerusalem Post, Netanyahu, Palestinian Authority, Jerusalem, Lebanon, Mubarak, Geneva, America, Richard Goldstone, Tunisia, Jeffrey Goldberg, Europe, Goldstone Report, UN, United States, Mahmoud Abbas, Cairo, East Jerusalem
--

Figure 4: Concepts identified by our mechanism for detecting contentious concepts in the CPAM data set.

## 5. Conclusion

In this paper, we described a computational model relied on the concept valuation to automatically discover and explore contentious concepts in two communities. The work is still in progress and we intend to extend the scale of CPAM onto very large volumes of social media, including instant messaging, micro blogs, as well as more official media. While CPAM work produced interesting insights into the structure of cultural narrative, the scale of the experiment was still too small for practical applications.

Another avenue of future work is to apply CPAM for influence strategy. Once the valuation change is accomplished for a contentious concept, the remaining elements of the valuation concept network are re-evaluated so that new "negotiable" items may be identified. The result is an influence strategy that could be used to define the order in which concepts can be negotiated. Here influence strategy means essentially "conflict reduction" – i.e., making the target group's valuation system more compatible with the influencer's valuation system.

A further extension work would be to account for socio-cultural functioning of the communities represented by the VCNs as 'living organisms'. We may consider a VCN as a primitive 'brain model' of the culture, in that it reflects the activation structure of important cultural concepts. Thus an open question would be whether VCNs can be extended to model how a culture reacts to, absorbs or rejects new information from the outside. This is different than the aforementioned influence strategy; rather, the focus here would be on the evolution of VCNs through adaptation to the external conditions.

## 6. References

- Breck, E.; Choi, Y.; and Cardie, C. 2007. Identifying expressions of opinion in context. In Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI).
- Brill, E. 1995. Transformation-based error-driven learning and natural language processing: A case study in part of speech tagging. *Computational Linguistics* 21:543–565.
- David, S., and Pinch, T. J. 2006. Six degrees of reputation: The use and abuse of online review and recommendation systems. *First Monday*. Special Issue on Commercial Applications of the Internet.
- Hu, M., and Liu, B. 2004. Mining opinion features in customer reviews. In Proceedings of AAAI, 755–760.
- Lager, T. 1999. The  $\mu$ -TBL System: Logic Programming Tools for Transformation-Based Learning. In Proceedings of the Third International Workshop on Computational Natural Language Learning (CoNLL'99).
- Small, S., Strzalkowski, T. and Webb, N. 2010. ECO: Effective Communication Online. Technical Report ILS-015, University at Albany, SUNY.
- Strapparava, C., and Mihalcea, R. 2008. Learning to Identify Emotions in Text. In Proceedings of the ACM Conference on Applied Computing ACM-SAC 2008.
- Toutanova, K., and Manning, C. D. 2000. Enriching the Knowledge Sources Used in a Maximum Entropy Part-of-Speech Tagger. In Proceedings of the Joint SIGDAT Conference on Empirical Methods in Natural Language Processing and Very Large Corpora (EMNLP/VLC-2000), 63–70.
- Vogt, T.; Andre', E.; and Bee, N. 2008. Emovoice – a framework for online recognition of emotions from voice. In PIT '08: Proceedings of the 4th IEEE tutorial and research workshop on Perception and Interactive Technologies for Speech-Based Systems, 188–199. Berlin, Heidelberg: Springer-Verlag.
- Webb, N.; Small, S.; and Shaikh, S. 2010. ECO Annotation Guidelines, Version 1.2. Technical Report ILS-012, University at Albany, SUNY.
- Wiebe, J.; Wilson, T.; and Cardie, C. 2005. Annotating expressions of opinions and emotions in language. *Journal of Language Resources and Evaluation* 39(2-3):165–210.